arXiv:1505.06794v1 [math.ST] 26 May 2015

Optimal Bayesian estimation in stochastic block models

Debdeep Pati

Department of Statistics, Florida State University, Tallahassee, FL,

email: debdeep@stat.fsu.edu

Anirban Bhattacharya

Department of Statistics, Texas A &M University, College Station, TX,

email: anirbanb@stat.tamu.edu

May 27, 2015

Abstract

With the advent of structured data in the form of social networks, genetic circuits and protein interaction networks, statistical analysis of networks has gained popularity over recent years. Stochastic block model constitutes a classical cluster-exhibiting random graph model for networks. There is a substantial amount of literature devoted to proposing strategies for estimating and inferring parameters of the model, both from classical and Bayesian viewpoints. Unlike the classical counterpart, there is however a dearth of theoretical results on the accuracy of estimation in the Bayesian setting. In this article, we undertake a theoretical investigation of the posterior distribution of the parameters in a stochastic block model. In particular, we show that one obtains optimal rates of posterior convergence with routinely used multinomial-Dirichlet priors on cluster indicators and uniform priors on the probabilities of the random edge indicators. En route, we develop geometric embedding techniques to exploit the lower dimensional structure of the parameter space which may be of independent interest.

KEYWORDS: Bayesian; block models; clustering; multinomial-Dirichlet; network; posterior contraction; random graph

1 Introduction

Data available in the form of networks are increasingly becoming common in applications ranging from brain connectivity, protein interactions, web applications and social networks to name a few, motivating an explosion of activity in the statistical analysis of networks in recent years Goldenberg et al. (2010). Estimating large networks offers unique challenges in terms of structured dimension reduction and estimation in stylized domains, necessitating new tools for inference. A rich variety of probabilistic models have been studied for network estimation, ranging from the classical Erdos and Renyi graphs Erdos & Rényi (1961), exponential random graph models Holland & Leinhardt (1981), stochastic block models Holland et al. (1983), markov graphs (Frank & Strauss, 1986) and latent space models (Hoff et al., 2002) to name a few.

In a network with n nodes, there are $O(n^2)$ possible connections betweens pairs of nodes, the exact number depending on whether the network is directed/undirected and whether self-loops are permitted. A common goal of the parametric models mentioned previously is to parsimoniously represent the $O(n^2)$ probabilities of connections between pairs of nodes in terms of fewer parameters. The stochastic block model achieves this by clustering the nodes into $k \ll n$ groups, with the probability of an edge between two nodes solely dependent on their cluster memberships. The block model originated in the mathematical sociology literature Holland et al. (1983), with subsequent widespread applications in statistics Wang & Wong (1987); Snijders & Nowicki (1997); Nowicki & Snijders (2001). In particular, the clustering property of block models offers a natural way to find *communities* within networks, inspiring a large literature on community detection Bickel & Chen (2009); Newman (2012); Zhao et al. (2012); Karrer & Newman (2011); Zhao et al. (2011); Amini et al. (2013). Various modifications of the stochastic block model have also been proposed, including the mixed membership stochastic block model Airoldi et al. (2009) and degree-corrected stochastic block model Zhao et al. (2012).

Statistical accuracy of parameter inference in the stochastic block model is of growing

interest, with one of the objects of interest being the $n \times n$ matrix of probabilities of edges between pairs of nodes, which we shall denote by $\theta = (\theta_{ij})$. Using a singular-value thresholding approach, Chatterjee (2014) obtained a $\sqrt{k/n}$ rate for estimating θ with respect to the squared ℓ_2 distance in a k-component stochastic block model. In a recent technical report, Gao et al. (2014) obtained an improved $k^2/n^2 + \log k/n$ rate by considering a least-squares type estimator. They also showed that the resulting rate is minimax-optimal; interestingly the minimax rate comprises of two parts which Gao et al. (2014) refer to as the *nonparametric* and *clustering* rates respectively. Among other related work, Bickel et al. (2013) provided conditions for asymptotic normality of maximum likelihood estimators in stochastic block models.

In this article, we consider a Bayesian formulation of a stochastic block model where θ is equipped with a hierarchical prior and study the convergence of the posterior distribution assuming the data to be generated from a stochastic block model. We show that one obtains the minimax rate of posterior convergence with essentially automatic prior choices, such as multinomial-Dirichlet priors on cluster indicators and uniform priors on the probability of the random edge indicators. Such priors are commonly used and there is a sizable literature Snijders & Nowicki (1997); Nowicki & Snijders (2001); Golightly & Wilkinson (2005); McDaid et al. (2013) on posterior sampling and inference in the stochastic block model. However, to the best of our knowledge, the present paper is the first to study the asymptotic properties of Bayesian estimation in stochastic block models.

Theoretical investigation of the posterior distribution in block models offers some unique challenges relative to the small but growing literature on posterior convergence in high-dimensional *sparse* problems Castillo & van der Vaart (2012); Pati et al. (2014); Banerjee & Ghosal (2014); Castillo et al. (2015). When a large subset of the parameters are exactly or approximately zero, the sparsity assumption can be exploited to reduce the complexity of the model space to derive tests for the true parameter versus the complement of a neighborhood of the true parameter Castillo & van der Vaart (2012); Pati et al.

(2014). It is now well appreciated that constructing such tests play a crucial role in posterior asymptotics Ghosal et al. (2000); Giné & Nickl (2011). In the present setting, we exploit the parsimonious structure of the parameters space as a result of clustering of n nodes into $k < n$ communities and develop geometric embedding techniques to derive such tests.

2 Preliminaries

For $\mathcal{S} \subset \mathbb{R}$, we shall denote the set of all $d \times d$ matrices with entries in $\mathcal{S}$ by $\mathcal{S}^{d\times d}$. For any $B = (B_{ll'}) \in \mathbb{R}^{d\times d}$, we denote the Euclidean (equivalently Frobenius) norm of B by $\|B\| = \sqrt{\sum_{l=1}^{d}\sum_{l'=1}^{d} B_{ll'}^2}$. Given $X^* \in \mathbb{R}^{d\times d}, W \in \mathbb{R}_+^{d\times d}$, let $\xi_{d^2}(X^*; W)$ denote the unit ellipsoid with *center* X^* and *weight* W given by

$$\xi_{d^2}(X^*; W) = \left\{ X \in \mathbb{R}^{d\times d} : \sum_{l=1}^{d}\sum_{l'=1}^{d} W_{ll'}(X_{ll'} - X^*_{ll'})^2 \leq 1 \right\}. \tag{1}$$

Viewed as a subset of $\mathbb{R}^{d^2}$, the Euclidean volume of $\xi_{d^2}(X^*; W)$, denoted by $|\xi_{d^2}(X^*; W)|$, is

$$|\xi_{d^2}(X^*; W)| = \frac{\pi^{d^2}}{\Gamma(d^2/2 + 1)} \prod_{l=1}^{d}\prod_{l'=1}^{d} W_{ll'}^{-1/2}. \tag{2}$$

Given sequences $\{a_n\}, \{b_n\}$, $a_n \lesssim b_n$ indicates there exists a constant $K > 0$ such that $a_n \leq K b_n$ for all large n. We say $a_n \asymp b_n$ when $a_n \lesssim b_n$ and $b_n \lesssim a_n$. Throughout, C, C' denote positive constants whose values might change from one line to the next.

3 Stochastic Block models

Let $A = (A_{ij}) \in \{0, 1\}^{n\times n}$ denote the adjacency matrix of a network with n nodes, with $A_{ij} = 1$ indicating the presence of an edge from node i to node j and $A_{ij} = 0$ indicating a lack thereof. To keep the subsequent notation clean, we shall consider directed networks with self-loops so that A_{ij} and A_{ji} need not be the same and A_{ii} can be both 0 and 1.

Our theoretical results can be trivially modified to undirected networks with or without self-loops.

Let θ_{ij} denote the probability of an edge from node i to j, with $A_{ij} \sim \text{Bernoulli}(\theta_{ij})$ independently for $1 \leq i, j \leq n$. A stochastic block model postulates that the nodes are clustered into communities, with the probability of an edge between two nodes solely dependent on their community memberships. Specifically, let $z_i \in \{1, \ldots, k\}$ denote the cluster membership of the ith node and $Q = (Q_{rs}) \in [0,1]^{k \times k}$ be a matrix of probabilities, with Q_{rs} indicating the probability of an edge from any node i in cluster r to any node j in cluster s. With these notations, a k-component stochastic block model is given by

$$A_{ij} \sim \text{Bernoulli}(\theta_{ij}), \quad \theta_{ij} = Q_{z_i z_j}. \tag{3}$$

We use $\mathbb{E}_\theta / \mathbb{P}_\theta$ to denote an expectation/probability under the sampling mechanism (3).

The stochastic block model clearly imposes a parsimonious structure on the node probabilities $\theta = (\theta_{ij})$ when $k \ll n$, reducing the effective number of parameters from $O(n^2)$ to $O(k^2 + n)$. To describe the parameter space for θ, we need to introduce some notations. For $k \leq n$, let $\mathcal{Z}_{n,k} = \big\{(z_1, \ldots, z_n) : z_i \in \{1, \ldots, k\}, 1 \leq i \leq n\big\}$ denote all possible clusterings of n nodes into k clusters. Elements of $\mathcal{Z}_{n,k}$ with be denoted by $z = (z_1, \ldots, z_n)$. For any $1 \leq r \leq k$, $z^{-1}(r)$ is used as a shorthand for $\{1 \leq i \leq n : z_i = r\}$; the nodes belonging to cluster r. When z is clear from the context, we shall use $n_r = |z^{-1}(r)|$ to denote the number of nodes in cluster r; clearly $\sum_{r=1}^k n_r = n$. With these notations, the parameter space Θ_k for θ is given by

$$\Theta_k = \{\theta \in [0,1]^{n \times n} : \theta_{ij} = Q_{z_i z_j},\ z \in \mathcal{Z}_{n,k},\, Q \in [0,1]^{k \times k}\}. \tag{4}$$

For any $z \in \mathcal{Z}_{n,k}$ and $Q \in [0,1]^{k \times k}$, we denote the corresponding $\theta \in \Theta_k$ by $\theta^{z,Q}$, so that $\theta^{z,Q}_{ij} = Q_{z_i z_j}$. In fact, $(z, Q) \mapsto \theta^{z,Q}$ is a surjective map from $\mathcal{Z}_{n,k} \times [0,1]^{k \times k} \to \Theta_k$, though it is clearly not injective.

Given $z \in \mathcal{Z}_{n,k}$, let $A_{[rs]}$ denote the $n_r \times n_s$ sub matrix of A consisting of entries A_{ij}

with $z_i = r$ and $z_j = s$. The joint likelihood of A under model (3) can be expressed as

$$P(A \mid z, Q) = \prod_{r=1}^{k} \prod_{s=1}^{k} P(A_{[rs]} \mid z, Q), \quad P(A_{[rs]} \mid z, Q) = \prod_{i:z_i=r} \prod_{j:z_j=s} Q_{rs}^{A_{ij}} (1 - Q_{rs})^{1-A_{ij}}. \tag{5}$$

A Bayesian specification of the stochastic block model can be completed by assigning independent priors to z and Q, which in turn induces a prior on Θ_k via the mapping $(z, Q) \mapsto \theta^{z,Q}$. We generically use $p(z, Q) = p(z)p(Q)$ to denote the joint prior on z and Q. The induced prior on Θ_k will be denoted by $\Pi(\theta)$ and the corresponding posterior given data $A = (A_{ij})$ will be denoted by $\Pi_n(\theta \mid A)$. The following fact is useful and heavily used in the sequel: for any $U \subset \Theta_k$,

$$\Pi(U) = \sum_{z \in \mathcal{Z}_{n,k}} \Pi(U \mid z)\ p(z) = \sum_{z \in \mathcal{Z}_{n,k}} p(Q : \theta^{z,Q} \in U)\ p(z), \tag{6}$$

where the second equality uses the independence of z and Q. Specific choices of $p(z)$ and $p(Q)$ are discussed below.

We assume independent $U(0, 1)$ prior on Q_{rs}. We consider a hierarchical prior on z where each node has probability π_r of being allocated to the rth cluster independently of the other nodes, and the vector of probabilities $\pi = (\pi_1, \ldots, \pi_k)$ follows a Dirichlet$(\alpha_1, \ldots, \alpha_k)$ prior. Here $\alpha_1, \ldots, \alpha_k$ are fixed hyper-parameters that do not depend on k or n; a default choice is $\alpha_r = 1/2$ for all $r = 1, \ldots, k$. Model (3) along with the prior specified above can be expressed hierarchically as follows:

$$Q_{rs} \overset{\text{ind}}{\sim} U(0, 1), \quad r, s = 1, \ldots, k, \tag{7}$$

$$P(z_i = k \mid \pi) = \pi_k, \quad i = 1, \ldots, n, \tag{8}$$

$$\pi \sim \text{Dirichlet}(\alpha_1, \ldots, \alpha_k), \tag{9}$$

$$A_{ij} \mid z, Q \overset{\text{ind}}{\sim} \text{Bernoulli}(\theta_{ij}), \quad \theta_{ij} = Q_{z_i z_j}. \tag{10}$$

A hierarchical specification as in (or very similar to) (7) – (10) has been commonly used in the literature; see for example, Snijders & Nowicki (1997); Nowicki & Snijders (2001); Golightly & Wilkinson (2005); McDaid et al. (2013). Analytic marginalizations

can be carried out due to the conjugate nature of the prior, facilitating posterior sampling McDaid et al. (2013). In particular, using standard multinomial-Dirichlet conjugacy, the marginal prior of z can be written as

$$p(z) = \frac{\Gamma(\sum_{r=1}^{k} \alpha_r)}{\Gamma(n + \sum_{r=1}^{k} \alpha_r)} \prod_{r=1}^{k} \frac{\Gamma(n_r + \alpha_r)}{\Gamma(\alpha_r)}, \quad z \in \mathcal{Z}_{n,k}, \tag{11}$$

where recall that $n_r = \sum_{i=1}^{n} \mathbb{1}(z_i = r)$. The following lemma provides a bound on the ratio of prior probabilities $p(z)/p(z')$ which is used subsequently in the proof of our main theorem.

Lemma 3.1. *Assume $z' \in \mathcal{Z}_{n,k}$ with $n'_r = \sum_{i=1}^{n} \mathbb{1}(z'_i = r) \geq 1$ for all $r = 1, \dots k$. Then, $\max_{z \in \mathcal{Z}_{n,k}} p(z)/p(z') \leq e^{Cn \log k}$, where C is a positive constant.*

Proof. Fix $z \in \mathcal{Z}_{n,k}$. From (11), $p(z)/p(z') = \prod_{r=1}^{k} \Gamma(n_r + \alpha_r)/\Gamma(n'_r + \alpha_r)$. Define non-negative integers $\beta_r = \lceil \alpha_r \rceil, \gamma_r = \lfloor \alpha_r \rfloor$, $\beta_\cdot = \sum_{r=1}^{k} \beta_r, \gamma_\cdot = \sum_{r=1}^{k} \gamma_r$. Recall the following facts about the Gamma function: (i) $\Gamma(x)$ is decreasing on $(0,1)$ with $1/(2x) \leq \Gamma(x) \leq 1/x$, (ii) $\Gamma(x) \leq 1$ for $x \in [1,2]$, $\Gamma(1) = \Gamma(2) = 1$ and (iii) $\Gamma(x)$ is increasing for $x \geq 2$. First, we claim $\Gamma(n'_r + \alpha_r) \geq C' \Gamma(n'_r + \gamma_r)$ for all r for an absolute constant $C' > 0$. To see this, recalling that $n'_r \geq 1$, separately consider the cases (a) $n'_r \geq 2$, (b) $n'_r = 1$ and $\alpha_r > 1$ and (c) $n'_r = 1$ and $\alpha_r < 1$. Cases (a) and (b) follow from fact (iii) above with $C' = 1$. For case (c), $\Gamma(n'_r + \alpha_r) = \Gamma(1 + \alpha_r) = \alpha_r \Gamma(\alpha_r) \geq 1/2$ by fact (i) and $\Gamma(n'_r + \gamma_r) = \Gamma(1) = 1$; therefore one may choose $C' = 1/2$ here. Next, we claim that $\Gamma(n_r + \alpha_r) \leq C \Gamma(n_r + \beta_r)$ for all r and an absolute constant $C > 0$. To see this, separately consider cases (a) $n_r \geq 1$, (b) $n_r = 0$ and $\alpha_r \geq 1$ and (c) $n_r = 0$ and $\alpha_r < 1$. Cases (a) and (b) once again follow from fact (iii) above with $C = 1$. For case (c), $\Gamma(n_r + \alpha_r) = \Gamma(\alpha_r)$ and $\Gamma(n_r + \beta_r) = \Gamma(1) = 1$, so one may choose $C = \max_{1 \leq r \leq k} \{1/\Gamma(\alpha_r)\}$. We thus have

$$\begin{aligned} \frac{p(z)}{p(z')} &= \prod_{r=1}^{k} \frac{\Gamma(n_r + \alpha_r)}{\Gamma(n'_r + \alpha_r)} \leq \left(\frac{C}{C'}\right)^k \prod_{r=1}^{k} \frac{\Gamma(n_r + \beta_r)}{\Gamma(n'_r + \gamma_r)} \\ &= \left(\frac{C}{C'}\right)^k \frac{(n + \beta_\cdot - k)!}{(n + \gamma_\cdot - k)!} \frac{\binom{n+\gamma_\cdot - k}{n'_1 + \gamma_1 - 1, \dots, n'_k + \gamma_k - 1}}{\binom{n+\beta_\cdot - k}{n_1 + \beta_1 - 1, \dots, n_k + \beta_k - 1}}, \end{aligned} \tag{12}$$

where we used $\sum_{r=1}^k n_r = \sum_{r=1}^k n'_r = n$ and $\binom{n}{m_1,\ldots,m_k} := n!/(m_1!\ldots m_k!)\mathbb{1}(m_1+\ldots m_k = n)$ is the multinomial coefficient. Since the multinomial coefficient as a function of $m_1,\ldots,m_{k-1}$, $(m_1,\ldots,m_{k-1}) \mapsto \binom{m}{m_1,\ldots,m_k}$, attains its minimum if $m_r = m$ for some r and $m_l = 0$ for $l \neq r$, and maximum if $m_r = m/k$ for all $r = 1,\ldots,k$, we have from (12),

$$\frac{p(z)}{p(z')} \leq \left(\frac{C}{C'}\right)^k \frac{(n+\beta_\cdot - k)!}{(n+\gamma_\cdot - k)!} \frac{(n+\gamma_\cdot - k)!}{[\{(n+\beta_\cdot - k)/k\}!]^k} = \left(\frac{C}{C'}\right)^k \frac{(n+\beta_\cdot - k)!}{[\{(n+\beta_\cdot - k)/k\}!]^k}.$$

Set $m = n + \gamma_\cdot - k$. Since $m! \asymp (m/e)^m\sqrt{m}$ by Stirling's bound, $m!/\{(m/k)!\}^k \lesssim k^m\sqrt{m}(k/m)^{k/2} \lesssim k^{Cn}$ for $C > 1$ large enough. The conclusion follows.

□

4 Posterior convergence rates in Stochastic Block Models

We are interested in concentration properties of the posterior $\Pi_n(\cdot \mid A)$ assuming the true data-generating parameter $\theta^0 \in \Theta_k$. To measure the discrepancy in the estimation of $\theta^0 \in \Theta_k$, the mean squared error has been used in the frequentist literature,

$$\frac{1}{n^2}\sum_{i=1}^n\sum_{j=1}^n(\hat{\theta}_{ij} - \theta^0_{ij})^2 = \frac{1}{n^2}\|\hat{\theta} - \theta^0\|^2, \tag{13}$$

where $\hat{\theta}$ is an estimator of θ^0. Chatterjee (2014) proposed estimating θ^0 using a low rank decomposition of the adjacency matrix A followed by a singular value decomposition to obtain a convergence rate of $\sqrt{k/n}$. More recently, Gao et al. (2014) considered a least squares type approach which can be related to maximum likelihood estimation where the Bernoulli likelihood is replaced by a Gaussian likelihood. They obtained a rate of $k^2/n^2 + \log k/n$, which they additionally showed to be the minimax rate over Θ_k, i.e.,

$$\inf_{\hat{\theta}} \sup_{\theta^0 \in \Theta_k} \mathbb{E}_{\theta_0}\frac{1}{n^2}\|\hat{\theta} - \theta^0\|^2 \asymp \frac{k^2}{n^2} + \frac{\log k}{n}. \tag{14}$$

Interestingly, the minimax rate has two components, k^2/n^2 and $\log k/n$. Gao et al. (2014) refer to the k^2/n^2 term in the minimax rate as the *nonparametric rate*, since it arises from

the need to estimate k^2 unknown elements in Q from n^2 observations. The second part, $\log k/n$, is termed as the *clustering rate*, which appears since the clustering configuration z is unknown and needs to be estimated from the data. Observe that the clustering rate grows logarithmically in k. Parameterizing $k = n^\delta$ with $\delta \in [0, 1]$, the interplay between the two components becomes clearer (refer to equation 2.6 of Gao et al. (2014)); in particular, the clustering rate dominates when k is small and the nonparametric rate dominates when k is large.

To evaluate Bayesian procedures from a frequentist standpoint, one seeks for the minimum possible sequence $\epsilon_n \to 0$ such that the posterior probability assigned to the complement of an ϵ_n-neighborhood (blown up by a constant factor) of θ^0 receives vanishingly small probabilities. The smallest such ϵ_n is called the *posterior convergence rate* Ghosal et al. (2000). There is now a growing body of literature showing that Bayesian procedures achieve the frequentist minimax rate of posterior contraction (up to a logarithmic term) in models where the parameter dimension grows with the sample size; see Bontemps (2011); Castillo & van der Vaart (2012); Pati et al. (2014); Banerjee & Ghosal (2014); van der Pas et al. (2014); Castillo et al. (2015) for some flavor of the recent literature.

We now state the main result of this article where we derive the convergence rate of the posterior arising from the hierarchical formulation (7) – (10).

Theorem 4.1. *Assume $A = (A_{ij})$ is generated from a k-component stochastic block model (3) with the true data-generating parameter $\theta^0 = (\theta^0_{ij}) \in \Theta_k$, where Θ_k is as in (4). Further assume that there exists a small constant $\delta \in (0, 1)$ such that $\theta^0_{ij} \in (\delta, 1 - \delta)$ for all $i, j = 1, \ldots, n$. Suppose the hierarchical Bayesian model (7) – (10) is fitted. Then, with $\epsilon_n^2 = k^2 \log(n/k)/n^2 + \log k/n$ and a sufficiently large constant $M > 0$,*

$$\lim_{n\to\infty} \mathbb{E}_{\theta^0} \Pi_n \left\{ \frac{1}{n^2} \left\| \theta - \theta^0 \right\|^2 > M^2 \epsilon_n^2 \mid A \right\} = 0. \tag{15}$$

Remark 4.1. *Since $\theta^0 \in \Theta_k$, following the discussion after (4), there exists $z^0 \in \mathcal{Z}_{n,k}$ and $Q^0 \in [0, 1]^{k\times k}$ such that $\theta^0 = \theta^{z^0, Q^0}$. The condition of the theorem posits that all entries*

of Q^0 lie in $(\delta, 1-\delta)$. The assumption $\theta^0 \in \Theta_k$ also implicitly implies that all the clusters have at least one observation, i.e., $n_r^0 = \sum_{i=1}^n \mathbb{1}(z_i^0 = r) \geq 1$ for all $r = 1, \ldots, k$; otherwise $\theta^0 \in \Theta_l$ for some $l < k$.

A proof of Theorem 4.1 can be found in Section 5. Theorem 4.1 shows that the posterior contracts at a (near) minimax rate of $k^2 \log(n/k)/n^2 + \log k/n$. The nonparametric component of the rate is slightly hurt by a logarithmic term; appearance of such an additional logarithmic term is common in Bayesian nonparametrics. An inspection of the proof additionally reveals that the techniques can be trivially extended to undirected graphs with or without self-loops and will produce the same rate.

4.1 Geometry of Θ_k

In this section, we derive a number of auxiliary results aimed at understanding the geometry of the parameter space Θ_k. These results are used to prove Theorem 4.1 and can be possibly of independent interest.

We first state a testing lemma which harnesses the ability of the likelihood to separate points in the parameter space.

Lemma 4.2. *Assume $\theta^0 \neq \theta^1 \in \Theta_k$ and let $E = \{\theta \in [0,1]^{n \times n} : \|\theta - \theta^0\| \leq \|\theta^1 - \theta^0\|/2\}$ be an Euclidean ball of radius $\|\theta^1 - \theta^0\|/2$ around θ^1 inside $[0,1]^{n \times n}$. Based on $A_{ij} \overset{ind}{\sim} Bernoulli(\theta_{ij})$ for $i, j = 1, \ldots, n$, consider testing $H_0 : \theta = \theta^0$ versus $H_1 : \theta \in E$. There exists a test function Φ such that*

$$\mathbb{E}_{\theta^0}(\Phi) \leq \exp\{-C_1 \|\theta^1 - \theta^0\|^2\}, \quad \sup_{\theta \in E} \mathbb{E}_\theta(1 - \Phi) \leq \exp\{-C_2 \|\theta^1 - \theta^0\|^2\}, \tag{16}$$

for constants $C_1, C_2 > 0$ independent of n, θ^1 and θ^0.

Proof. Define the test function Φ as

$$\Phi = \mathbb{1}\Big\{ \sum_{i=1}^n \sum_{j=1}^n (\theta_{ij}^1 - \theta_{ij}^0)(A_{ij} - \theta_{ij}^0) > \|\theta^1 - \theta^0\|^2/4 \Big\},$$

where $\mathbb{1}(\cdot)$ denotes the indicator of a set. We show below that this test has the desired error rates (16).

We first bound the type-I error $\mathbb{E}_{\theta^0}(\Phi)$. Noting that under $\mathbb{P}_{\theta^0}$, $(A_{ij} - \theta^0_{ij})$ are independent zero mean random variables with $|A_{ij} - \theta^0_{ij}| < 1$, we use a version of Hoeffding's inequality (refer to Proposition 5.10 of Vershynin (2010)) to conclude that,

$$\begin{aligned}\mathbb{E}_{\theta^0}(\Phi) &= \mathbb{P}_{\theta^0}\Big\{ \sum_{i=1}^n \sum_{j=1}^n (\theta^1_{ij} - \theta^0_{ij})(A_{ij} - \theta^0_{ij}) > \big\|\theta^1 - \theta^0\big\|^2/4 \Big\} \\ &\le \exp\Big\{ - C_1 \frac{\big\|\theta^1 - \theta^0\big\|^4}{\big\|\theta^1 - \theta^0\big\|^2} \Big\} = \exp\big\{ - C_1 \big\|\theta^1 - \theta^0\big\|^2 \big\}\end{aligned}$$

for a constant $C_1 > 0$ independent of n, θ^1 and θ^0.

We next bound the type-II error $\sup_{\theta \in E} \mathbb{E}_\theta(1 - \Phi)$. Fix $\theta \in E$. We have,

$$\begin{aligned}\mathbb{E}_\theta(1 - \Phi) &= \mathbb{P}_\theta\Big\{ \sum_{i=1}^n \sum_{j=1}^n (\theta^1_{ij} - \theta^0_{ij})(A_{ij} - \theta^0_{ij}) < \big\|\theta^1 - \theta^0\big\|^2/4 \Big\} \\ &= \mathbb{P}_\theta\Big\{ \sum_{i=1}^n \sum_{j=1}^n (\theta^1_{ij} - \theta^0_{ij})(A_{ij} - \theta_{ij}) < \big\|\theta^1 - \theta^0\big\|^2/4 - \big\langle \theta^1 - \theta^0, \theta - \theta^0 \big\rangle \Big\}, \quad (17)\end{aligned}$$

where we abbreviate $\langle \theta', \theta'' \rangle = \sum_{i=1}^n \sum_{j=1}^n \theta'_{ij} \theta''_{ij}$. Bound

$$\begin{aligned}&\big\langle \theta^1 - \theta^0, \theta - \theta^0 \big\rangle \\ &= \big\langle \theta^1 - \theta^0, \theta^1 - \theta^0 \big\rangle - \big\langle \theta^1 - \theta^0, \theta^1 - \theta \big\rangle \\ &\ge \big\|\theta^1 - \theta^0\big\|^2 - \big\|\theta^1 - \theta^0\big\|^2/2 = \big\|\theta^1 - \theta^0\big\|^2/2,\end{aligned}$$

where the penultimate step used the Cauchy–Schwartz inequality along with the fact that $\big\|\theta - \theta^1\big\| \le \big\|\theta^1 - \theta^0\big\|/2$. Substituting in (17) and noting that under $\mathbb{P}_\theta$, $(A_{ij} - \theta_{ij})$ are independent zero mean bounded random variables, another application of Hoeffding's inequality yields

$$\begin{aligned}\mathbb{E}_\theta(1 - \Phi) &\le \mathbb{P}_\theta\Big\{ \sum_{i=1}^n \sum_{j=1}^n (\theta^1_{ij} - \theta^0_{ij})(A_{ij} - \theta_{ij}) < -\big\|\theta^1 - \theta^0\big\|^2/4 \Big\} \\ &\le \exp\Big\{ - C_2 \frac{\big\|\theta^1 - \theta^0\big\|^4}{\big\|\theta^1 - \theta^0\big\|^2} \Big\} = \exp\big\{ - C_2 \big\|\theta^1 - \theta^0\big\|^2 \big\}\end{aligned}$$

for some constant $C_2 > 0$ independent of n and θ. Taking a supremum over $\theta \in E$ yields the desired result. $\square$

Our next result is concerned with the structure of a specific type of Euclidean balls inside Θ_k. Recall that $\theta^{z,Q}$ denotes the element of Θ_k with $\theta^{z,Q}_{ij} = Q_{z_i z_j}$. For $z \in \mathcal{Z}_{n,k}$, let

$$\Theta_k(z) = \left\{\theta^{z,Q} : Q \in [0,1]^{k\times k}\right\} \tag{18}$$

denote a slice of Θ_k along z. In other words, given z, $\Theta_k(z)$ is the image of the map $Q \mapsto \theta^{z,Q}$ in Θ_k. Suppose $\theta^* = \theta^{z^*,Q^*} \in \Theta_k$, and consider a ball $B(z)$ in $\Theta_k(z)$ centered at θ^* of the form $B(z) = \left\{\theta \in \Theta_k(z) : \left\|\theta - \theta^*\right\| < t\right\}$ for some $t > 0$. If $z^* = z$, then it is straightforward to observe that

$$\left\|\theta^{z,Q} - \theta^{z^*,Q^*}\right\|^2 = \sum_{r=1}^{k}\sum_{s=1}^{k} n_r n_s (Q_{rs} - Q^*_{rs})^2, \tag{19}$$

where recall that $n_r = \sum_{i=1}^{n} \mathbb{1}(z_i = r)$ for $r = 1, \ldots, k$. Therefore, although a subset of $[0,1]^{n\times n}$, $B(z)$ can be identified with a k^2-dimensional ellipsoid in $[0,1]^{k\times k}$. When $z^* \neq z$, one no longer has a nice identity as above and the geometry of $B(z)$ is more difficult to describe. However, we show below in Lemma 4.3 that $B(z)$ is always contained inside a set $\widetilde{B}(z)$ in $\Theta_k(z)$ which can be identified with a k^2-dimensional ellipsoid in $[0,1]^{k\times k}$. Recall our convention for describing ellipsoids from (1).

Lemma 4.3. *Fix $z^* \in \mathcal{Z}_{n,k}, Q^* \in [0,1]^{k\times k}$, and let $\theta^* = \theta^{z^*,Q^*}$. For $z \in \mathcal{Z}_{n,k}$ and $t > 0$, let $B(z) = \left\{\theta \in \Theta_k(z) : \left\|\theta - \theta^*\right\| < t\right\}$. Set $W_{rs} = n_r n_s / t^2$ and $W = (W_{rs})$, where $n_r = \sum_{i=1}^{n} \mathbb{1}(z_i = r)$ for $r = 1, \ldots, k$. Then, $B(z) \subseteq \widetilde{B}(z)$, where*

$$\widetilde{B}(z) = \left\{\theta^{z,Q} : Q \in \xi_{k^2}(\bar{Q}^*, W) \cap [0,1]^{k\times k}\right\} \tag{20}$$

for some $\bar{Q}^* \in [0,1]^{k\times k}$ depending on Q^*, z^* and z. In particular, if $z^* = z$, then $\bar{Q}^* = Q^*$ and the containment becomes equality, i.e., $B(z) = \widetilde{B}(z)$.

Remark 4.2. *From* (1), $\xi_{k^2}(\bar{Q}^*, W)$ *in Lemma 4.3 is the collection of all* Q *satisfying* $\sum_{r=1}^k \sum_{s=1}^k n_r n_s (Q_{rs} - \bar{Q}^*_{rs})^2 < t^2$. *The last part of Lemma 4.3 is consistent with the discussion preceding* (19). *When* $z^* = z$, (19) *implies that* $B(z)$ *consists of all* $\theta^{z,Q}$ *with* $Q \in [0,1]^{k\times k}$ *satisfying* $\sum_{r=1}^k \sum_{s=1}^k n_r n_s (Q_{rs} - Q^*_{rs})^2 < t^2$.

Proof. We begin by constructing $\bar{Q}^*$. For $1 \le r, r' \le k$, let $I_{r,r'} = z^{-1}(r) \cap (z^*)^{-1}(r')$ and $n_{r,r'} = |I_{r,r'}|$. Clearly $\{I_{r,r'}, r' = 1, \ldots, k\}$ is a partition of $z^{-1}(r)$ and hence $n_r = \sum_{r'=1}^k n_{r,r'}$. With these notations, define

$$\bar{Q}^*_{rs} = \frac{1}{n_r n_s} \sum_{r'=1}^k \sum_{s'=1}^k n_{r,r'} n_{s,s'} Q^*_{r',s'}. \tag{21}$$

Clearly, $\bar{Q}^*_{rs}$ is a weighted average of $Q^*_{r',s'}$ with weights proportional to $n_{r,r'} n_{s,s'}$ and therefore $\bar{Q}^* \in [0,1]^{k\times k}$. For $\theta^{z,Q} \in \Theta_k(z)$, we have

$$\left\|\theta^{z,Q} - \theta^{z^*,Q^*}\right\|^2 = \sum_{i=1}^n \sum_{j=1}^n \left(Q_{z_i z_j} - Q^*_{z^*_i z^*_j}\right)^2.$$

Expanding the squares, the term $\sum_{i=1}^n \sum_{j=1}^n Q^2_{z_i z_j} = \sum_{r=1}^k \sum_{s=1}^k n_r n_s Q^2_{rs}$. The cross product term can be simplified to

$$\begin{aligned}\sum_{i=1}^n \sum_{j=1}^n Q_{z_i z_j} Q^*_{z^*_i z^*_j} &= \sum_{r=1}^k \sum_{s=1}^k \sum_{r'=1}^k \sum_{s'=1}^k n_{r,r'} n_{s,s'} Q_{rs} Q_{r's'} \\ &= \sum_{r=1}^k \sum_{s=1}^k n_r n_s Q_{rs} \bar{Q}^*_{rs}. \end{aligned} \tag{22}$$

In view of these identities, we can write

$$\begin{aligned}&\sum_{i=1}^n \sum_{j=1}^n \left(Q_{z_i z_j} - Q^*_{z^*_i z^*_j}\right)^2 = \sum_{r=1}^k \sum_{s=1}^k n_r n_s (Q_{rs} - \bar{Q}^*_{rs})^2 + \\ &\sum_{i=1}^n \sum_{j=1}^n (Q^*_{z^*_i z^*_j})^2 - \sum_{r=1}^k \sum_{s=1}^k n_r n_s (\bar{Q}^*_{rs})^2. \end{aligned} \tag{23}$$

We shall show below that the expression in (23) is non-negative. This completes the proof, since we then have $\sum_{r=1}^k \sum_{s=1}^k n_r n_s (Q_{rs} - \bar{Q}^*_{rs})^2 \le \left\|\theta^{z,Q} - \theta^{z^*,Q^*}\right\|^2 < t^2$ (see Remark 4.2).

Recalling the definition of $\bar{Q}^*_{rs}$ from (21), the expression in (23) can be written as

$$\sum_{i=1}^{n}\sum_{j=1}^{n}(Q^*_{z_i^* z_j^*})^2 - \sum_{r=1}^{k}\sum_{s=1}^{k} n_r n_s (\bar{Q}^*_{rs})^2$$
$$= \sum_{r=1}^{k}\sum_{s=1}^{k}\Big[\sum_{r'=1}^{k}\sum_{s'=1}^{k} n_{r,r'} n_{s,s'} (Q^*_{r's'})^2 - \frac{1}{n_r n_s}\Big\{\sum_{r'=1}^{k}\sum_{s'=1}^{k} n_{r,r'} n_{s,s'} Q^*_{r's'}\Big\}^2\Big].$$

The non-negativity of this quantity now follows from the Cauchy–Schwartz inequality $(\sum_\alpha a_\alpha b_\alpha)^2 \le \sum_\alpha a_\alpha^2 b_\alpha \sum_\alpha b_\alpha$, with $a_\alpha = Q^*_{r's'}$ and $b_\alpha = n_{r,r'} n_{s,s'}$; and the fact that $\sum_{r'=1}^{k} n_{r,r'} = n_r$.

When $z^* = z$, it is clear that $n_{r,r'} = \delta_{rr'}$ and hence $\bar{Q}^* = Q^*$. All subsequent inequalities then become equalities and the last part is proved.

□

Corollary 4.4. *Inspecting the proof of Lemma 4.3, the condition $Q \in [0,1]^{k\times k}$ is only used to show that barQ$^* \in [0,1]^{k\times k}$. If we let Q to be unrestricted, then the containment relation continues to hold as subsets of $\mathbb{R}^{k\times k}$, i.e.,*

$$\Big\{\theta^{z,Q} : Q \in \mathbb{R}^{k\times k}, \big\|\theta^{z,Q} - \theta^{z^*,Q^*}\big\| < t\Big\} \subseteq \Big\{\theta^{z,Q} : Q \in \xi_{k^2}(\bar{Q}^*, W)\Big\}, \tag{24}$$

with equality when $z^* = z$.

Lemma 4.3 and Corollary 4.4 crucially exploit the lower dimensional structure underlying the parameter space Θ_k and is used subsequently multiple times. First, recall from (6) that one needs a handle on $p(Q : \theta^{z,Q} \in U)$ to bound the prior probability of $U \subset \Theta_k$. In particular, if $U = \{\big\|\theta - \theta^0\big\| < t\}$, then $p(Q : \theta^{z,Q} \in U)$ equals the volume of $U \cap \Theta_k(z)$, which can be suitably bounded by the volume of the bounding k^2 dimensional ellipsoid. Second, a handle on the size of balls in Θ_k facilitates calculating the complexity of the model space (in terms of metric entropy) which is pivotal in proving the posterior concentration; in particular, to extend the test function in Lemma 4.2 to construct test functions against more complex alternatives in Lemma 4.5 below. Once again, the dimensionality reduction is key to preventing the metric entropy from growing too fast.

Lemma 4.5. *Recall ϵ_n from Theorem 4.1. Assume $\theta^0 \in \Theta_k$ and for $l \geq 1$, let $U_{l,n} = \{\theta \in \Theta_k : ln\epsilon_n \leq \|\theta - \theta^0\| < (l+1)n\epsilon_n\}$. Based on $A_{ij} \overset{ind}{\sim} Bernoulli(\theta_{ij})$ for $i, j = 1, \ldots, n$, consider testing $H_0 : \theta = \theta^0$ versus $H_1 : \theta \in U_{l,n}$. There exists a test function $\Phi_{l,n}$ such that*

$$\mathbb{E}_{\theta^0}(\Phi_{l,n}) \leq \exp(-C_1 l^2 n^2 \epsilon_n^2), \quad \sup_{\theta \in U_{l,n}} \mathbb{E}_\theta(1 - \Phi_{l,n}) \leq \exp(-C_2 l^2 n^2 \epsilon_n^2), \tag{25}$$

for constants $C_1, C_2 > 0$ independent of n.

Proof. Since $\theta^0 \in \Theta_k$, there exists $z^0 \in \mathcal{Z}_{n,k}$ and $Q^0 \in [0,1]^{k \times k}$ with $\theta^0 = \theta^{z^0, Q^0}$. For $z \in \mathcal{Z}_{n,k}$, define $U_{l,n}(z) = U_{l,n} \cap \Theta_k(z)$, where $\Theta_k(z)$ is as in (18). Clearly,

$$U_{l,n}(z) = \{\theta^{z,Q} : Q \in [0,1]^{k \times k},\ ln\epsilon_n \leq \|\theta^{z,Q} - \theta^{z^0,Q^0}\| < (l+1)n\epsilon_n\}, \tag{26}$$

and $U_{l,n} \subset \cup_{z \in \mathcal{Z}_{n,k}} U_{l,n}(z)$. We first use Lemma 4.2 to construct tests against $U_{l,n}(z)$ for fixed z. Our desired test is obtained by taking the maximum of all such test functions.

Fix $z \in \mathcal{Z}_{n,k}$. Let $\mathcal{N}_{l,n}(z) = \{\theta_{l,n,h} \in U_{l,n}(z) : h \in I_{l,n}(z)\}$ be a *maximal* $ln\epsilon_n/2$-separated set inside $U_{l,n}(z)$ for some index set $I_{l,n}(z)$; i.e., $\mathcal{N}_{l,n}(z)$ is such that $\|\theta^1 - \theta^2\| \geq ln\epsilon_n/2$ for all $\theta^1 \neq \theta^2 \in \mathcal{N}_{l,n}(z)$, and no subset of $U_{l,n}(z)$ containing $\mathcal{N}_{l,n}(z)$ has this property. We provide a volume argument to determine an upper bound for $|I_{l,n}(z)|$, the cardinality of $\mathcal{N}_{l,n}(z)$. The separation property implies that Euclidean balls of radius $ln\epsilon_n/4$ centered at the points in $\mathcal{N}_{l,n}(z)$ are disjoint. Since $B_h^+ := \{\theta^{z,Q} : Q \in \mathbb{R}^{k \times k}, \|\theta^{z,Q} - \theta_{l,n,h}\| < ln\epsilon_n/4\}$ is contained inside an Euclidean ball of radius $ln\epsilon_n/4$ centered at $\theta_{l,n,h}$, the sets B_h^+ are disjoint as h varies over $I_{l,n}(z)$. By the triangle inequality, all B_h^+s lie inside $B^+ = \{\theta^{z,Q} : Q \in \mathbb{R}^{k \times k}, \|\theta^{z,Q} - \theta^0\| \leq (5l/4 + 1)n\epsilon_n\}$, since $\|\theta^{z,Q} - \theta^0\| \leq \|\theta^{z,Q} - \theta_{l,n,h}\| + \|\theta_{l,n,h} - \theta^0\| \leq (l+1)n\epsilon_n + ln\epsilon_n/4$.

It should be noted that the sets B_h^+s and B^+ are constructed in a way that Q is not restricted to be inside $[0,1]^{k \times k}$. This allows us to invoke Corollary 4.4 to identify B_h^+ and B^+ with appropriate ellipsoids in $\mathbb{R}^{k^2}$ and simplify volume calculations. First, since $\theta_{l,n,h} \in \Theta_k(z)$ for each h, it follows from (the equality part of) Corollary 4.4 that $B_h^+ = \{\theta^{z,Q} : Q \in \xi_{k^2}(\bar{Q}_h, \widetilde{W})\}$ with $\bar{Q}_h$ constructed as in the proof of Lemma 4.3 and $\widetilde{W}_{rs} =$

$n_r n_s/\{(ln\epsilon_n)^2\}$. The equality is crucially used below; also note that $\widetilde{W}$ does not depend on h. Invoking Corollary 4.4 one more time, we obtain $B^+ \subset \{\theta^{z,Q} : Q \in \xi_{k^2}(\bar{Q}^0, W)\}$, with $W_{rs} = n_r n_s/[\{(5l/4+1)n\epsilon_n\}^2]$. We conclude that the Euclidean ellipsoids $\xi_{k^2}(\bar{Q}_h, \widetilde{W})$ are disjoint as h varies over $I_{l,n}(z)$ and all of them are contained in $\xi_{k^2}(\bar{Q}^0, W)$. Comparing volumes,

$$|\xi_{k^2}(\bar{Q}_h, \widetilde{W})||I_{l,n}(z)| \leq |\xi_{k^2}(\bar{Q}^0, W)|.$$

Using the volume formula in (2) and canceling out common terms, we finally have

$$|I_{l,n}(z)| \leq \left\{\frac{(5l/4+1)}{l/2}\right\}^{k^2} \leq 9^{k^2}. \tag{27}$$

We are now in a position to construct the test. The maximality of $\mathcal{N}_{l,n}(z)$ implies that $\mathcal{N}_{l,n}(z)$ is an $ln\epsilon_n/2$-net of $U_{l,n}(z)$, i.e., the sets $E_{l,n,z,h} = \{\theta \in [0,1]^{n\times n} : \|\theta - \theta_{l,n,h}\| < ln\epsilon_n/2\}$ cover $U_{l,n}(z)$ as h varies. For each $\theta_{l,n,h} \in \mathcal{N}_{l,n}(z)$, consider testing $H_0 : \theta = \theta^0$ versus $H_1 : \theta \in E_{l,n,z,h}$ using the test function from Lemma 4.2. Lemma 4.2 is applicable since $\|\theta^0 - \theta_{l,n,h}\| \geq ln\epsilon_n$; let $\Phi_{l,n,z,h}$ denote the corresponding test with type-I and II errors bounded above by $e^{-Cl^2n^2\epsilon_n^2}$. Define $\Phi_{l,n} = \max_{z\in\mathcal{Z}_{n,k}} \max_{h\in I_{l,n}(z)} \Phi_{l,n,z,h}$. For any $\theta \in U_{l,n}$, there exists $z \in \mathcal{Z}_{n,k}$ and $h \in I_{l,n}(z)$ such that $\theta \in E_{l,n,z,h}$, so that $\mathbb{E}_\theta(1-\Phi_{l,n}) \leq \mathbb{E}_\theta(1-\Phi_{l,n,z,h}) \leq e^{-Cl^2n^2\epsilon_n^2}$. Taking supremum over $\theta \in U_{l,n}$ delivers the desired type-II error. Further, the type-I error of $\Phi_{l,n}$ can be bounded as

$$\mathbb{E}_{\theta^0}(\Phi_{l,n}) \leq \sum_{z\in\mathcal{Z}_{n,k}} \sum_{h\in I_{l,n}(z)} \mathbb{E}_{\theta^0}(\Phi_{l,n,z,h}) \leq k^n 9^{k^2} e^{-Cl^2n^2\epsilon_n^2}, \tag{28}$$

since $|\mathcal{Z}_{n,k}| = k^n$ and by (27), $|I_{l,n}(z)| \leq 9^{k^2}$ for all z. The conclusion follows since $n^2\epsilon_n^2 = k^2\log(n/k) + n\log k \gtrsim k^2 + n\log k$.

□

5 Proof of Theorem 4.1

Proof. Let $\mathbb{E}_0/\mathbb{P}_0$ denote an abbreviation to $\mathbb{E}_{\theta^0}/\mathbb{P}_{\theta^0}$. Since $\theta^0 \in \Theta_k$, there exists $z^0 \in \mathcal{Z}_{n,k}$ and $Q^0 \in [0,1]^{k\times k}$ with $\theta^0 = \theta^{z^0,Q^0}$. Recall $\epsilon_n = k^2\log n/n^2 + \log k/n$ and define

$U_n = \{\theta \in \Theta_k : \|\theta - \theta^0\|^2 > M^2 n^2 \epsilon_n^2\}$ for some large constant $M > 0$ to be chosen later. Letting $f_{\theta_{ij}}(A_{ij}) = \theta_{ij}^{A_{ij}}(1-\theta_{ij})^{A_{ij}}$ denote the Bernoulli(θ_{ij}) likelihood evaluated at A_{ij}, the posterior probability assigned to U_n can be written as

$$\Pi_n(U_n \mid A) = \frac{\int_{U_n} \prod_{i=1}^n \prod_{j=1}^n \frac{f_{\theta_{ij}}(A_{ij})}{f_{\theta_{ij}^0}(A_{ij})}\, p(dz, dQ)}{\int_{\Theta_k} \prod_{i=1}^n \prod_{j=1}^n \frac{f_{\theta_{ij}}(A_{ij})}{f_{\theta_{ij}^0}(A_{ij})}\, p(dz, dQ)} = \frac{\mathcal{N}_n}{\mathcal{D}_n}, \tag{29}$$

where $\mathcal{N}_n$ and $\mathcal{D}_n$ respectively denote the numerator and denominator of the fraction in (29). Let $\mathcal{F}_n$ denote the σ-field generated by $\tilde{A} = (\tilde{A}_{ij})$, with $\tilde{A}_{ij}$ independently distributed as Bernoulli(θ_{ij}^0); the true data generating distribution. We first claim that there exists a set $\mathcal{A}_n \in \mathcal{F}_n$ where we can bound $\mathcal{D}_n$ from below with large probability under $\mathbb{P}_0$ in Lemma 5.1. A proof can be adapted from Lemma 10 of Ghosal & van der Vaart (2007) and hence omitted.

Lemma 5.1. *Assume θ^0 satisfies the conditions of Theorem 4.1. Then, there exists a set $\mathcal{A}_n$ in the σ-field $\mathcal{F}_n$ with $\lim_{n\to\infty} \mathbb{P}_0(\mathcal{A}_n) = 1$ such that within $\mathcal{A}_n$,*

$$\mathcal{D}_n \geq e^{-Cn^2\epsilon_n^2} \Pi\big(\|\theta - \theta^0\|^2 < n^2\epsilon_n^2\big).$$

In view of Lemma 5.1, it is sufficient to prove that

$$\lim_{n\to\infty} \mathbb{E}_0\big\{\Pi_n(U_n \mid A)\mathbb{1}_{\mathcal{A}_n^c}\big\} = 0.$$

For $l \geq M$, let $U_{l,n} = \{\theta \in \Theta_k : l^2n^2\epsilon_n^2 \leq \|\theta - \theta^0\|^2 < (l+1)^2n^2\epsilon_n^2\}$ denote an annulus in Θ_k centered at θ^0 with inner and outer Euclidean radii $ln\epsilon_n$ and $(l+1)n\epsilon_n$ respectively. Using a standard testing argument (see, for example, the proof of Proposition 5.1 in Castillo & van der Vaart (2012)) in conjunction with Lemma 5.1, one arrives at

$$\mathbb{E}_0\big\{\Pi_n(U_n \mid A)\mathbb{1}_{\mathcal{A}_n^c}\big\} \leq \sum_{l=M}^{\infty} \left\{\mathbb{E}_0(\Phi_{l,n}) + \beta_{l,n} \sup_{\theta \in U_{l,n}} \mathbb{E}_\theta(1 - \Phi_{l,n})\right\} \tag{30}$$

where

$$\beta_{l,n} = \frac{\Pi(U_{l,n})}{e^{-Cn^2\epsilon_n^2}\Pi(\|\theta - \theta^0\|^2 < n^2\epsilon_n^2)} \tag{31}$$

and $\Phi_{l,n}$ is the test function constructed in Lemma 4.5 for testing $H_0 : \theta = \theta^0$ versus $H_1 : \theta \in U_{l,n}$ with error rates as in (25). Recall $U_{l,n}(z) = U_{l,n} \cap \Theta_k(z)$ and its equivalent representation in (26) from the proof of Lemma 4.5. Since $U_{l,n} \subseteq \cup_{z \in \mathcal{Z}_{n,k}} U_{l,n}(z)$, from (6),

$$\Pi(U_{l,n}) \leq \sum_{z \in \mathcal{Z}_{n,k}} \Pi\big\{U_{l,n}(z)\big\} = \sum_{z \in \mathcal{Z}_{n,k}} \Pi\big\{U_{l,n}(z) \mid z\big\}\, p(z), \tag{32}$$

where $p(z)$ is the prior probability (11) of z under the Dirichlet-multinomial prior. By an application of Lemma 4.3, $U_{l,n}(z) \subset \{\theta^{z,Q} : Q \in \xi_{k^2}(\bar{Q}^0, W) \cap [0,1]^{k \times k}\}$ with $W_{rs} = n_r n_s / \{(l+1) n \epsilon_n\}^2, 1 \leq r, s \leq k$. Therefore, $\Pi\{U_{l,n}(z) \mid z\}$ is bounded above by the probability of the set $\xi_{k^2}(\bar{Q}^0, W) \cap [0,1]^{k \times k}$ under the uniform prior on Q, which in turn can be bounded above by the Euclidean volume of $\xi_{k^2}(\bar{Q}^0, W)$. Using volume formula (2),

$$\Pi\big\{U_{l,n}(z) \mid z\big\} \leq |\xi_{k^2}(\bar{Q}^0, W)| = \frac{\pi^{k^2}}{\Gamma(k^2/2+1)} \prod_{r=1}^{k} \prod_{s=1}^{k} \frac{(l+1) n \epsilon_n}{n_r n_s}. \tag{33}$$

Next, consider the term $\Pi(\|\theta - \theta^0\|^2 < n^2 \epsilon_n^2)$ in the denominator of the expression for $\beta_{l,n}$. Bound $\Pi(\|\theta - \theta^0\|^2 < n^2 \epsilon_n^2) \geq \Pi(\|\theta - \theta^0\|^2 < n^2 \epsilon_n^2 \mid z = z^0) p(z^0)$ and using Lemma 4.3 once again,

$$\Pi\Big(\|\theta - \theta^0\|^2 < n^2 \epsilon_n^2 \mid z = z^0\Big) = p\Big\{Q : \sum_{r=1}^{k} \sum_{s=1}^{k} n_{0r} n_{0s} (Q_{rs} - Q^0_{rs})^2 < n^2 \epsilon_n^2\Big\}. \tag{34}$$

The probability in the right hand side of the above display is the volume of the *intersection* of an ellipsoid with $[0,1]^{k \times k}$, and therefore we cannot simply replace the probability by the volume of the ellipsoid. Instead, we embed an appropriate rectangle inside the intersection of the ellipsoid and $[0,1]^{k \times k}$. We claim that

$$\prod_{r=1}^{k} \prod_{s=1}^{k} [Q^0_{rs} - \epsilon_n/2, Q^0_{rs} + \epsilon_n/2] \subset \Big\{Q \in [0,1]^{k \times k} : \sum_{r=1}^{k} \sum_{s=1}^{k} n_{0r} n_{0s} (Q_{rs} - Q^0_{rs})^2 < n^2 \epsilon_n^2\Big\}. \tag{35}$$

First, based on our assumption that all entries of Q^0 are bounded away from 0 and 1 and the fact that $\epsilon_n \to 0$, it is immediate that the rectangle is contained in $[0,1]^{k \times k}$ for sufficiently large n. Second, for any Q with $|Q_{rs} - Q^0_{rs}| \leq \epsilon_n/2$ for all $1 \leq r, s \leq k$, we have

$$\sum_{r=1}^{k} \sum_{s=1}^{k} n_{0r} n_{0s} (Q_{rs} - Q^0_{rs})^2 \leq \frac{\epsilon_n^2}{4} \sum_{r=1}^{k} \sum_{s=1}^{k} n_{0r} n_{0s} = \frac{n^2 \epsilon_n^2}{4},$$

thereby proving the claim in (35). Now we can bound $\Pi(\left\|\theta - \theta^0\right\|^2 < n^2\epsilon_n^2 \mid z = z^0)$ from below by the volume of the rectangle, which equals $\epsilon_n^{k^2}$. Using this fact along with the bounds (32), (33) and (34), we have from (31) that

$$e^{-Cn^2\epsilon_n^2}\beta_{l,n} \leq \frac{\pi^{k^2}(l+1)^{k^2}n^{k^2}}{\Gamma(k^2/2+1)} \sum_{z \in \mathcal{Z}_{n,k}} \frac{p(z)}{p(z^0)}. \tag{36}$$

Since $n_{0r} \geq 1$ for all $r = 1, \ldots, k$, invoke Lemma 3.1 to bound $\sum_{z \in \mathcal{Z}_{n,k}} p(z)/p(z_0) \leq |\mathcal{Z}_{n,k}|e^{Cn\log k} \leq e^{Cn\log k}$, since $|\mathcal{Z}_{n,k}| = k^n$. Next, use the well-known fact (see, for example, Abramowitz & Stegun (1964)) that for any $\alpha > 0$, $\Gamma(\alpha+1) \geq \sqrt{2\pi}e^{-\alpha}\alpha^{\alpha+1/2}$ to obtain

$$e^{-Cn^2\epsilon_n^2}\beta_{l,n} \lesssim \{\pi\sqrt{2e}(l+1)\}^{k^2}\left(\frac{n}{k}\right)^{k^2} \lesssim \{\pi\sqrt{2e}(l+1)\}^{k^2}e^{k^2\log(n/k)}. \tag{37}$$

Substituting in (30), the expression in (30) converges to zero for all M larger than a suitable constant.

□

6 Discussion

In this article, we presented a theoretical investigation of posterior contraction in stochastic block models. One crucial assumption in our current results is that the true number of clusters k is known. An interesting direction is to develop a fully Bayesian approach by placing a prior on k and to show that the corresponding procedure yields optimal rates of posterior convergence adaptively for all values of $k \in \{1, 2, \ldots, n\}$. Such an approach can be connected to nonparametric estimation of networks Bickel & Chen (2009) where one typically assumes a more flexible way of data generation; $A_{ij} \mid \xi_i, \xi_j \sim \text{Bernoulli}\{f(\xi_i, \xi_j)\}$, where f is a function from $[0,1]^2 \to [0,1]$, called a *graphon* and ξ_is are iid random variables on $[0,1]$. It is well known (refer, for example, to Airoldi et al. (2013)) that one can approximate a sufficiently smooth graphon using elements of Θ_k. When the smoothness of the graphon is unknown, the prior on k should facilitate the posterior to concentrate in

the appropriate region. Using such approximation results and modifying our Theorem 4.1, it may be possible to derive posterior convergence rates for estimating a graphon.

References

Abramowitz, M. & Stegun, I. (1964). *Handbook of mathematical functions: with formulas, graphs, and mathematical tables*. No. 55. Courier Corporation.

Airoldi, E., Costa, T. & Chan, S. (2013). Stochastic blockmodel approximation of a graphon: Theory and consistent estimation. In *Advances in Neural Information Processing Systems*.

Airoldi, E. M., Blei, D. M., Fienberg, S. E. & Xing, E. P. (2009). Mixed membership stochastic blockmodels. In *Advances in Neural Information Processing Systems*.

Amini, A. A., Chen, A., Bickel, P. J. & Levina, E. (2013). Pseudo-likelihood methods for community detection in large sparse networks. *The Annals of Statistics* **41**, 2097–2122.

Banerjee, S. & Ghosal, S. (2014). Posterior convergence rates for estimating large precision matrices using graphical models. *Electronic Journal of Statistics* **8**, 2111–2137.

Bickel, P. & Chen, A. (2009). A nonparametric view of network models and Newman–Girvan and other modularities. *Proceedings of the National Academy of Sciences* **106**, 21068–21073.

Bickel, P. J., Choi, D., Chang, X. & Zhang, H. (2013). Asymptotic normality of maximum likelihood and its variational approximation for stochastic blockmodels. *The Annals of Statistics* **41**, 1922–1943.

Bontemps, D. (2011). Bernstein–von Mises theorems for Gaussian regression with increasing number of regressors. *The Annals of Statistics* **39**, 2557–2584.

Castillo, I., Schmidt-Hieber, J. & van der Vaart, A. W. (2015). Bayesian linear regression with sparse priors. *The Annals of Statistics* (to appear).

Castillo, I. & van der Vaart, A. W. (2012). Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics* **40**, 2069–2101.

Chatterjee, S. (2014). Matrix estimation by universal singular value thresholding. *The Annals of Statistics* **43**, 177–214.

Erdos, P. & Rényi, A. (1961). On the evolution of random graphs. *Bull. Inst. Internat. Statist* **38**, 343–347.

Frank, O. & Strauss, D. (1986). Markov graphs. *Journal of the American Statistical association* **81**, 832–842.

Gao, C., Lu, Y. & Zhou, H. H. (2014). Rate-optimal graphon estimation. *arXiv preprint arXiv:1410.5837* .

Ghosal, S., Ghosh, J. K. & van der Vaart, A. W. (2000). Convergence rates of posterior distributions. *Annals of Statistics* **28**, 500–531.

Ghosal, S. & van der Vaart, A. W. (2007). Convergence rates of posterior distributions for noniid observations. *The Annals of Statistics* **35**, 192–223.

Giné, E. & Nickl, R. (2011). Rates of contraction for posterior distributions in l_r-metrics, $1 \leq r \leq \infty$. *The Annals of Statistics* **39**, 2883–2911.

Goldenberg, A., Zheng, A., Fienberg, S. & Airoldi, E. (2010). A survey of statistical network models. *Foundations and Trends® in Machine Learning* **2**, 129–233.

Golightly, A. & Wilkinson, D. J. (2005). Bayesian inference for stochastic kinetic models using a diffusion approximation. *Biometrics* **61**, 781–788.

Hoff, P. D., Raftery, A. E. & Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical association* **97**, 1090–1098.

Holland, P. W., Laskey, K. B. & Leinhardt, S. (1983). Stochastic blockmodels: First steps. *Social networks* **5**, 109–137.

Holland, P. W. & Leinhardt, S. (1981). An exponential family of probability distributions for directed graphs. *Journal of the American Statistical association* **76**, 33–50.

Karrer, B. & Newman, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Physical Review E* **83**, 016107.

McDaid, A., Murphy, T. B., Friel, N. & Hurley, N. (2013). Improved bayesian inference for the stochastic block model with application to large networks. *Computational Statistics & Data Analysis* **60**, 12–31.

Newman, M. E. J. (2012). Communities, modules and large-scale structure in networks. *Nature Physics* **8**, 25–31.

Nowicki, K. & Snijders, T. A. B. (2001). Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association* **96**, 1077–1087.

Pati, D., Bhattacharya, A., Pillai, N. S. & Dunson, D. (2014). Posterior contraction in sparse bayesian factor models for massive covariance matrices. *The Annals of Statistics* **42**, 1102–1130.

Snijders, T. A. B. & Nowicki, K. (1997). Estimation and prediction for stochastic blockmodels for graphs with latent block structure. *Journal of classification* **14**, 75–100.

van der Pas, S., Kleijn, B. & van der Vaart, A. (2014). The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics* **8**, 2585–2618.

Vershynin, R. (2010). Introduction to the non-asymptotic analysis of random matrices. *Arxiv preprint arxiv:1011.3027* .

Wang, Y. J. & Wong, G. Y. (1987). Stochastic blockmodels for directed graphs. *Journal of the American Statistical Association* **82**, 8–19.

Zhao, Y., Levina, E. & Zhu, J. (2011). Community extraction for social networks. *Proceedings of the National Academy of Sciences* **108**, 7321–7326.

Zhao, Y., Levina, E. & Zhu, J. (2012). Consistency of community detection in networks under degree-corrected stochastic block models. *The Annals of Statistics* **40**, 2266–2292.